Scanpath Theory in Virtual Reality: Spatiotemporal Aspects of Eye and Head Movement
Similarity Predict Memory Performance

Nicola C. Anderson¹

Oliver Jacobs¹

Walter F. Bischof¹

Alan Kingstone¹

¹Department of Psychology, University of British Columbia

Correspondence should be addressed to Nicola C Anderson, Department of Psychology,
University of British Columbia, 2136 West Mall, Vancouver, BC, V6T 1Z4
Email: nicola.anderson@ubc.ca

Abstract

It has long been thought that visual perception is represented in sensorimotor processes that unfold over time. One prominent theory predicts that our memory for a scene consists of both the scene content and the motor commands (i.e., eye movements) used to explore that scene. This Scanpath Theory (Noton & Stark, Science 171 (1971) 308-311) has long been contested, with many studies providing evidence both for, and against it. That past work, however, has failed to account for the fact that visual perception is embodied within an active system of effectors, namely, that people routinely move both their eyes and head to explore visible space. In the present work we tested Scanpath Theory while observers were free to move within a 360-degree VR environment. Their task was to encode and later recognise panoramic scenes within this fully immersive world. During both encoding and recognition, we recorded their eye and head movements using a VR headset equipped with eye and head tracking. Our results reveal that eye and head movement patterns are diagnostic of memory performance; and that scene recognition improves when certain movements that had occurred during encoding are repeated. Finally, including head movement measures enhances performance prediction, strengthening the evidence for Scanpath Theory, and reinforcing the fact that the head moves in service of the eyes in allocating attention.

Keywords

Virtual Reality; Eye and head tracking; Visual Attention; Scanpath Theory; Recurrence Quantification Analysis

1. Introduction

“In summary Noton’s interesting but rather informally presented theory fails to accommodate all pertinent phenomena, and as yet is not supported by conclusive evidence. Nonetheless, Noton has advanced an original and nontrivial idea. In the final analysis such may be a researcher’s only true *raison d’être*.”

-Allen D. Allen on “A theory of visual pattern perception” by
David Noton (Allen & Noton, 1971, p. 389)

Scanpath Theory, introduced by Noton (1970), and the seminal papers by Noton and Stark (1971a, 1971b, 1971c), encompass the idea that memory for an image or pattern consists of both the visual features that a person looks at, and the sequence of eye movements that link these features together. This “scanpath” is idiosyncratic for both the person and the image, and explicitly consists of “an alternating sequence of sensory and motor components” (Noton, 1970, p. 357). The more contentious part of this theory is the suggestion that the eye movements used to look at something are not only for perception, but are an “integral part of the memories on which recognition is based” (Noton & Stark, 1971b). Thus, eye movements are functional for memory, and it is not only important where a person looks, but in what order.

The early comment on Noton’s (1970) theory in the above quote (Allen & Noton, 1971) foreshadowed decades of research on his idea, highlighting just how compelling his theory is. One general approach to studying Scanpath Theory has been to compare the sequence of eye movements made when studying a set of images (‘encoding’), to the sequence of eye movements made when either imagining or recognizing the same images (e.g., Brandt & Stark, 1997;

Foulsham & Underwood, 2008). Such research has sparked an explosion of interest in scanpath comparison techniques (e.g., Anderson et al., 2015; Cristino et al., 2010; Dewhurst et al., 2012; Fahimi & Bruce, 2020; Le Meur & Baccino, 2013; Mannan et al., 1995). It has generally been observed that scanpaths tend to be the most similar for the same observer viewing the same image (Anderson et al., 2015; Foulsham et al., 2012) corroborating the initial observations of Noton and Stark, and other early eye movement researchers (e.g., Buswell, 1935; Yarbus, 1967).

In addition to understanding the relative contributions of observers and image features to scanpath variability, a key prediction in Noton and Stark's theory is that scanpaths should be functional for memory (Noton, 1970; Noton & Stark, 1971b). Many researchers in both visual imagery and memory have found evidence for a functional role of eye movements (Althoff & Cohen, 1999; Brandt & Stark, 1997; Foulsham & Underwood, 2008; Harding & Bloj, 2010; Holm & Mäntylä, 2007; Johansson et al., 2011; Kolers, 1973; Laeng & Teodorescu, 2002; Tremblay et al., 2006; Tulving & Thomson, 1973; Whiteside, 1978; Zelinsky & Loschky, 2005) and against it (Foulsham et al., 2012; Godijn & Theeuwes, 2012; Hollingworth & Henderson, 2002; Humphrey & Underwood, 2010; Locher & Nodine, 1974; Richardson et al., 2009; Richardson & Spivey, 2000). These contradictory findings are likely due in part to the fact that eye movements, or scanpaths made when encoding an image, do not always clearly predict performance, and when they do, the evidence is correlational in nature (c.f., Foulsham & Kingstone, 2013), and may only be predictive for limited aspects of a scanpath, such as fixation duration similarity (Foulsham et al., 2012). When tested experimentally, by constraining fixations during encoding or recognition, it has been shown that re-presenting the image locations that were fixated during encoding improves recognition performance, but a strict replay of the order of these re-presentations does not (Foulsham & Kingstone, 2013). More recent experimental evidence suggests that the extent to

which observers explore a scene during encoding predicts memory performance, and that eye movements during recognition serve to disambiguate encoded scenes, from new, but similar, scenes (Damiano & Walther, 2019). Taken together, it seems that while eye movements during encoding may predict memory performance, it is not necessary for the order of these movements to be reinstated for successful image recognition.

While it appears that Noton and Stark's Scanpath Theory may be too inflexible to account for a functional role for scanpaths in memory performance, the fundamental idea that the experience of looking at a natural scene is composed of both the features that one looks at, as well as the motor commands used to look there, holds true. Noton (1970) thought that the motor commands used to shift attention between features were a key component of his theory, but also specified that his theory would hold similarly for shifts of attention without saccades ("internal attention-shifting activity" p. 350). One can also turn this possibility on its head (both figuratively and literally) and suppose that the motor commands used to shift attention between features may also include shifts beyond saccades, to include shifts of the head. The present paper is concerned precisely with this possibility.

Until recently almost all of the research on Scanpath Theory, as well as the methods developed in studying it, have been concerned with shifts of the eyes alone, with participants required to keep their head still. There is growing recognition in the field, however, that eye movement behaviour collected in the lab when an observer's head is discouraged to move, is often not representative of eye movement behaviour in everyday life (Backhaus et al., 2020; Hooge et al., 2019; Risko et al., 2016; 't Hart et al., 2009). For example, when people walk about their environment they normally turn their head to look at things that interest them, the result being that their eyes tend to remain relatively central with respect to the head (Foulsham et al., 2011;

Foulsham & Kingstone, 2017). Similarly, when observers are free to rotate their head in line with rotated text, they do so, to offset the cognitive load of mentally rotating the text (Risko et al., 2014). Indeed, there is a growing recognition that there is a close link between head movements and cognitive processing, especially memory (Li et al., 2016, 2018; Rajsic et al., 2018; e.g., Võ & Wolfe, 2015). For example, Solman and Kingstone (2014) demonstrated that memory improves for items that are looked at when eye movements are accompanied by a head movement, presumably because it takes more energy to move the head and the eyes than the eyes alone (i.e., having worked harder to look at something people encode it more deeply into memory).

Scanpath theory fits well within a field that is increasingly responsive to the idea that much of our cognition is embodied (and action-based), and embedded within the natural world. Historically, scanpath comparison research has been limited to traditional screen-based tracking, where only the eyes are allowed to move, and the energetic cost of orienting is low. In addition, an observer's field of view has typically been restricted to the size of a display monitor (usually between 40 – 50 degrees visual angle), which means that all the relevant visual information is delivered simultaneously to an observer without them having to move their head. Given the recent work indicating that natural viewing behaviour is predominantly head-centered (Foulsham et al., 2011), and memory performance improves when an eye movement is accompanied by a head movement (Solman & Kingstone, 2014) it follows that new evidence that favours Scanpath Theory may be found when observers are free to move their head and eyes to take in visual information of their choosing.

Noton and Stark were limited with respect to the technological tools available to test their theory, in a similar way that for decades later researchers routinely needed to constrain the head and stimulus size. Fortunately, recent advances in Virtual Reality (VR) have provided the ability

to track simultaneously both the eyes and head with respect to a common reference frame (see Hessels et al., 2018; Hooge et al., 2019). In the present work, we test Scanpath Theory in a context where participants are free to move their eyes and head to explore 360 degree panoramic scenes. Participants were instructed to freely view a series of 360 degree panoramas, after which they performed a memory test where half of the images had been seen before, and half were new. We ask 1) What kind of spatio-temporal gaze and head movement behaviours lead to better recognition performance? And, 2) Does the addition of head movement measures lead to better modeling and prediction of memory performance?

To answer these questions, we performed several analyses. We examined the extent to which standard measures of head and eye movements predicted recognition performance (e.g., fixation count, fixation duration, etc.). In addition, we used cross recurrence analysis (CRA; Anderson et al., 2015; Dale et al., 2011) to quantify the overlap between the temporal components of head and eye movements to examine which specific behaviours lead to better recognition memory. Finally, we compare the ability of a combined gaze and head scanpath model to predict performance, over and above gaze and head movement models alone. In sum, we tested Scanpath Theory by assessing the spatial and temporal overlap of unconstrained head and eye movements, both at encoding and recognition, within fully immersive panoramic visual scenes.

2. Methods

2.1. Participants

Twenty-one undergraduate participants (15 female) from the University of British Columbia participated in this experiment for course credit. All reported normal or corrected-to-normal vision. Participants provided informed consent prior to participation, and the study was approved by the ethics board of the University of British Columbia (H10-00527).

2.2. Apparatus

This study was conducted on a custom-built desktop computer running an HTC Vive Virtual Reality device equipped with an SMI eye tracker. The HTC Vive headset has a resolution of 1080x1200 pixels, with a 110° by 113° field of view and a refresh rate of 90 Hz. The virtual experimental set-up was built using the Unity platform, a popular game development software (*Unity Technologies*, 2017). The SMI eye tracker has a sampling rate of 250 Hz and was controlled using the SMI-designed Unity plugin. Tracking accuracy was maintained by performing a calibration every 20 trials. Calibration consisted of following a moving white circle with a red dot in the middle to 5 different regions. The SMI Unity plugin reports only success or failure of calibration, not accuracy, so calibration was repeated upon failure. Head movements (movements of the HTC Vive headset) were tracked via two infrared base stations (60 Hz) located on opposite corners of the room. Participants were provided with brief instructions on handling the HTC Vive headset, were seated in a swivel chair in the center of the room (in clear view of both base stations), were handed the keyboard, and directed to the response key (Keypad Enter).

2.3. Stimuli

The virtual space for both the encoding and recognition phases consisted of a sphere around the participant, onto which different indoor and outdoor omni-directional panoramic scenes were projected, such that the participant appeared to be immersed in the scenes. The scenes taken from the SUN360 Panorama Database (Xiao et al., 2010) and had a resolution of 4096 x 2160 pixels.

2.4. Procedure

Participants were instructed that there would be two phases to the study, an image encoding phase and an image recognition phase.

In the encoding phase, participants were familiarized with the VR equipment and given the instructions to remember each of the pictures they were about to see for a later memory test. Each trial began with a uniform gray sphere with a fixation cross that allowed for the experimenter to monitor tracker accuracy. Participants pressed the ‘Spacebar’ key on a keyboard that they were holding to indicate that they were ready for the trial to begin. Each picture was presented for 10 seconds. There were 80 trials in the encoding phase with a random combination of 40 indoor and 40 outdoor scenes. This part of the experiment took approximately 20 minutes, after which participants were instructed to remove the headset and take a 5-minute break.

In the recognition phase, each trial began with a uniform gray sphere with a fixation cross. Participants were again immersed in the VR environment and pressed the ‘Spacebar’ key to indicate that they were ready for the trial to begin. Each trial involved the 10 second presentation of one of 80 scenes, with 40 indoor and 40 outdoor scenes, half of each type were old (i.e., drawn from the original encoding stimulus set), and half were new. After each image was presented, instructions on a gray background asked participants to press the far left ‘Ctrl’ key if they had seen the image before, and the far right keypad ‘Enter’ key if they had not seen the image before.

3. Analysis and Results

3.1. Data handling

Using the SMI Unity plugin, and treating gaze direction as a 3D vector, we recorded each participant's gaze direction in space at 250 Hz. Head position (position of the headset) was recorded at the fastest possible rate (which in Unity, is ~ 60 Hz). Using the head direction vector, which was linearly interpolated to 250 Hz, and gaze direction vector, gaze position relative to the virtual sphere was extracted. Note that from our data recordings, actual gaze sampling occurred at an average rate of 232.95 Hz. In this way the 3D gaze data were converted into positions on the virtual sphere which could then be examined in a manner equivalent to traditional monitor-based studies.

To clarify the reference frame being dealt with (Hessels et al., 2018), we tracked gaze and head with respect to the common reference frame of the sphere. That is, gaze and head position are referenced to the virtual sphere, which is equivalent to world-based coordinates. Most lab-based eye tracking studies measure eye movements on a screen, where the frame of reference is world-centered. In contrast, mobile eye tracking tends to track gaze position with respect to the world camera, where the frame of reference moves with the head (i.e., head-centered, or eye-in-head).

The gaze or head positions were transformed into longitude and latitude values on the virtual sphere, with longitudes in the range $[-180, 180]$ degrees (west to east) and latitudes in the range $[-90, 90]$ degrees (south pole to north pole). Note that the initial fixation cross and the center of the panoramic images was at coordinates $[0, 0]$. As the SMI outputs only sample data, it was necessary to perform event detection on the raw gaze data. Fixations were defined as stable gaze points and

were extracted using the IDT algorithm (Salvucci & Goldberg, 2000; Blignaut, 2009; Komogortsev et al., 2010) using a minimum duration of 80 ms and a maximum dispersion threshold of 3° angular distance (Bischof et al., 2020). Head movements are not ballistic in the way that eye movements are. For this reason, there are, in contrast to gaze, no natural demarcations for head shifts and head fixations. In order to compare head with gaze, we adopted the convention that, for every gaze fixation, we define a head fixation as the average head hit point on the sphere during that fixation.

The temporal analysis reported below uses fixation and head fixation data in the form of scanpaths or ‘headpaths,’ which are the successive spherical positions over time in a trial. An example scan- and head-path is shown in Figure 1.

Figure 1: Example scanpath (green circles and lines) and headpath (blue squares and lines) subset from the encoding trial of a random participant. Note that the image in this example is semi-

transparent (which was not the case in the actual experiment), and has been zoomed in around the scanpath and headpath for clarity.

All statistical analysis and follow-up tests were conducted using the R packages lmer (Bates et al., 2015), sjPlot (Lüdtke, 2020), afex (Singmann et al., 2019), emmeans (Lenth, 2019), and pROC (Robin et al., 2011); figures were produced using ggplot2 (Wickham, 2016). Fixations with durations longer than 800 ms were removed from the analysis (1.7% of data removed).

3.2. Generalized Linear Mixed Effects Models

In order to account for the fact that participants vary in terms of their ability to remember scenes, and for the fact that some images may be more or less memorable than others, we used generalized linear mixed effects models (GLMM's) to assess performance (correct vs. incorrect). We used the lme4 package (Bates et al., 2015), with participants and images as random effects (allowing only the intercept to vary), with each of the gaze and head measures as fixed effects. As many of the measures had different scales, they were mean scaled before modeling. Models used the default optimizer ('bobyqa'), and where non-convergence occurred, the number of iterations was increased to 200,000.

We performed separate GLMM's for each measure, to determine whether it predicted performance (on its own) over and above any variance accounted for by the random effects of image and participants. Likelihood ratio tests were conducted using the 'mixed' function from the afex package (Singmann et al., 2019), and for model estimates, approximated p-values from Wald z-statistics were used. For ease of interpretation, log odds output from the GLMM's were converted to odds ratios using the sjPlot package (Lüdtke, 2020).

3.3. Recognition Performance

Recognition accuracy was good, with 95% ($SD = 3.11\%$) accuracy for correctly recognizing indoor scenes, and 88% ($SD = 4.45\%$) accuracy for correctly recognizing outdoor scenes. Outdoor scenes were significantly harder to recognize, $t(20) = 4.27$, $p < .001$.

3.4. General eye movement measures

We examined the extent to which general eye and head movement measures during encoding and recognition predicted memory performance. In Table 1 for encoding, we report eye and head movements to only those images that were later included in the recognition phase. Note, however, there were no differences in any of the eye and head movement measures between those images that were included in the recognition phase and those that were not (all t 's < 1.88 , all p 's > 0.21). One subject was removed from the analysis of the encoding data for having a perfect score, while eight subjects were removed from the analysis of the recognition data for missing at least one category of responses (e.g., one subject made no false alarms).

In order to determine whether a given measure predicted memory performance, each measure was submitted to a GLMM, separately for encoding and recognition. Including participants and images as random effects, this GLMM takes into account a given person's ability to remember images, as well as a given image's memorability (details in Section Generalized Linear Mixed Effects Models).

Table 1: Means and standard deviations of general eye movement measures. For encoding trials, only those images that were later included in the recognition phase were included.

	Encoding	Old	New
Number of fixations per trial			
Correct (Hit/Correct Rejection)	39.27 (3.39)	32.31 (8.91)	36.53 (6.82)
Incorrect (Miss/False Alarm)	38.14 (5.09)	38.46 (6.65)	39.06 (6.83)
Mean fixation duration (ms)			
Correct (Hit/Correct Rejection)	196.3 (19.3)	216.6 (24.1)	203.7 (18.0)
Incorrect (Miss/False Alarm)	199.2 (31.7)	192.6 (17.4)	189.9 (19.1)
Gaze central tendency (degrees visual angle)			
Correct (Hit/Correct Rejection)	59.82 (10.46)	37.70 (15.9)	45.15 (15.2)
Incorrect (Miss/False Alarm)	57.14 (12.02)	50.48 (15.6)	54.21(20.3)
Head central tendency (degrees visual angle)			
Correct (Hit/Correct Rejection)	49.72 (11.27)	30.40 (7.97)	36.75 (9.47)
Incorrect (Miss/False Alarm)	47.35 (12.39)	42.35 (11.6)	44.75 (18.1)

Table 1 shows the total number of fixations for correct and incorrect trials in the encoding and recognition phases. In the encoding phase, there were a greater number of fixations in images that were correctly remembered, compared to those that were not, $\chi^2(1) = 10.66$, $p = .001$, *Odds Ratio* = 1.37, *CI* = [1.14, 1.65], $z = 3.36$, $p < .001$. Note that due to singularity issues arising from near-zero variance for the random effect of subject, it was dropped from the model, following Barr and colleagues (2013). In the recognition phase, there were significantly fewer fixations for old images that were correctly recognized (hits), compared to those that were not (misses), $\chi^2(1) = 49.36$, $p < .001$, *Odds Ratio* = 0.33, *CI* = [0.23, 0.47], $z = -5.99$, $p < .001$. For new images, there were significantly fewer fixations for correct rejections, compared to false alarms, $\chi^2(1) = 5.22$, $p = .022$, *Odds Ratio* = 0.60, *CI* = [0.38, 0.95], $z = -2.16$, $p = .031$.

Table 1 also shows mean fixation durations for correct and incorrect trials in the encoding and recognition phase. In the encoding phase, fixation durations were marginally shorter for correctly remembered images, compared to those that were not, $\chi^2(1) = 3.60, p = .058$, *Odds Ratio* = 0.83, *CI* = [0.69, 1.00], $z = -1.99, p = .047$. In the recognition phase, fixation durations were significantly longer for old images that were correctly recognized, compared to those that were not, $\chi^2(1) = 36.83, p < .001$, *Odds Ratio* = 3.05, *CI* = [1.97, 4.71], $z = 5.03, p < .001$. For new images, fixation durations were significantly longer for correct rejections, compared to false alarms, $\chi^2(1) = 4.75, p = .029$, *Odds Ratio* = 1.81, *CI* = [1.04, 3.16], $z = 2.08, p = .037$.

Table 1 also shows the mean central tendency of gaze for correct and incorrect trials in the encoding and recognition phase. Central tendency (in degrees visual angle) was calculated by taking the great circle distance between each fixation position and the location of the initial fixation. In the encoding phase, fixations were less centralized in images that were correctly recognized, compared to those that were not, $\chi^2(1) = 7.37, p = .007$, *Odds Ratio* = 1.33, *CI* = [1.08, 1.63], $z = 2.74, p = .006$. Note that due to singularity issues arising from near-zero variance for the random effect of subject, it was dropped from the model (Barr et al., (2013)). In the recognition phase, fixations were more central for old images that were correctly recognized compared to those that were not, $\chi^2(1) = 53.60, p < .001$, *Odds Ratio* = 0.38, *CI* = [0.28, 0.51], $z = -6.54, p < .001$. For new images, fixations were more central for correct rejections, compared to false alarms, $\chi^2(1) = 6.24, p = .012$, *Odds Ratio* = 0.97, *CI* = [0.94, 0.99], $z = -2.51, p = .012$.

Finally, Table 1 shows the mean central tendency of the head for correct and incorrect trials in the encoding and recognition phase. In the encoding phase, head positions were less centralized in images that were later correctly recognized, compared to those that were not, $\chi^2(1) = 3.98, p = .046$, *Odds Ratio* = 1.04, *CI* = [1.00, 1.08], $z = 2.00, p = .046$. In the recognition phase, head

positions were more central for old images that were correctly recognized compared to those that were not , $\chi^2(1) = 341.14, p < .001$, *Odds Ratio* = 0.71, *CI* = [0.69, 0.74], $z = -18.55, p < .001$. For new images, head positions were more central for correct rejections, compared to false alarms, $\chi^2(1) = 18.07, p < .001$, *Odds Ratio* = 0.88, *CI* = [0.83, 0.93], $z = -4.26, p < .001$.

In summary, the extent to which individuals move their gaze and head to explore panoramic scenes is associated with their memory performance. During encoding, more fixations, and more exploration (with both the eyes and head), was associated with scenes that were later correctly recognized. During recognition, correctly recognized images were associated with fewer fixations, but longer fixation durations and less exploration with the eyes and head. There was a similar pattern for correct rejections compared to false alarms, but the magnitude of these differences is less pronounced (see Table 1).

3.5. Scanpath Similarity

In the previous section we found that the extent to which people move their eyes and head to explore a scene is predictive of their ability to recognize it later. In this section, we examine if there is something about the particular ways in which people explore the scenes that may be diagnostic of their memory performance. To achieve this we compared encoding scanpaths to recognition scanpaths using a technique called cross-recurrence analysis (Bischof et al., 2019; Dale et al., 2011), which in turn is a generalization of recurrence quantification analysis (Anderson et al., 2013). This technique allows one to assess the extent to which recognition fixations match encoding fixations over both space and time, and to quantify the specific aspects of eye and head movement behaviour that predicts successful recognition of panoramic scenes.

An example of a cross-recurrence pattern is illustrated in Figure 2, Panel D. Consider two fixation sequences $\mathbf{E}$ (for encoding) and $\mathbf{R}$ (for recognition) that are adjusted to have the same lengths. Within these sequences, two fixations $\mathbf{e}_i$ and $\mathbf{r}_j$ are cross-recurrent if they match, i.e., their distance is below a particular threshold, in the present case 10° angular distance. CRA was conducted on the fixation sequences (scanpaths) and head position sequences (headpaths) of a given participant encoding and then recognizing each of the 40 images that were repeated during the recognition phase (a within-subject, within-image comparison). From the cross recurrence map (Figure 2, Panel D), several key measures can be computed that quantify aspects of scanpath (and headpath) similarity that are directly interpretable. Below is a brief note on the choice of radius size, followed by a description of the eye and head movement behaviour that each measure quantifies, but more formal definitions of each measure are given in the Appendix.

3.6. Choice of radius size

The number and percentage of recurrences is directly related to the choice of the radius size. With increasing radii, recurrence increases. Following Webber and Zbilut (2005), we sought to keep the percentage of overall recurrence fairly small (~7% for gaze), and by keeping the radius the same of gaze and head, we could more directly compare the two. Nevertheless, it is worth noting that more than half (~54%) of fixations in the recognition trials were recurrent with at least one encoding fixation.

3.7. CRA measures

Recurrence is the percentage of red dots in the recurrence plot (out of the total possible recurrences) and represents the amount of spatial overlap between encoding and recognition fixations and head positions (see Figure 2D).

Determinism is the percentage of these recurrent points that form diagonal lines in the cross recurrence plot. This particular measure most closely resembles more ‘pure’ scanpath comparison measures, as it quantifies the number of fixations or head positions that overlap in both space and time (i.e., fixation trajectories that are common to both fixation sequences – see Panel B, bright green boxes, in Figure 2). The *maxline* measure captures the length of the longest diagonal line in a recurrence plot. This value would be higher when a greater sequential number of encoding and recognition fixations or head positions overlapped.

Vertical Laminarity and the closely related *Vertical Trapping Time* quantify the percentage of recurrent points that form vertical lines in the plot (see Panel C and pink boxes in Figure 2). Vertical laminarity is a measure of repeated fixations or head positions on a particular region that are common to both scanpath or headpath sequences. For example, it corresponds to a single fixation in the encoding scanpath that is recurrent with a sequence of fixations in the recognition

scanpath. This captures the detailed inspection during image recognition of a region that was looked at during encoding. Vertical Trapping Time is the average length of these vertical lines, indicating the average sequential number of fixations or head positions that tend to land on previously encoded regions, i.e., how ‘trapped’ recognition fixations or head positions are on previously encoded regions.

Horizontal Laminarity represents the percentage of recurrent points that form horizontal lines in the plot. This captures the behaviour where an individual may scan particular encoding regions in detail, only briefly looking at them during recognition (see the last few (33-35) encoding fixations in Figure 2). *Horizontal Trapping Time* is the average length of these horizontal lines, i.e., how ‘trapped’ encoding fixations or head positions were on regions that were covered during recognition. These measures capture the idea that particular encoding regions that garnered much attention are revisited during recognition.

The *Center of Recurrence Mass* (CORM) measure quantifies where on the cross recurrence plot the majority of recurrences are. A large positive value indicates that during recognition, observers spend more time at locations looked at earlier during encoding (a primacy effect). A large negative corm value indicates that during recognition, observers spend more time at locations looked at later during encoding (a recency effect).

The *Clusters* measure quantifies the percentage of recurrent fixations that form clusters with more than 8 fixations within an 8-connected neighbourhood (“Pixel Connectivity,” 2021) in the recurrence plot (in Panel D of Figure 2 there are 3 clusters, one at the bottom left with 8 fixations, one in the center with 46 fixations, and one on the right with 25 fixations— these result from clusters

of overlapping encoding and recognition fixations in the scene center, near the high dresser, and the window in Panel A).

Figure 2: Panel A: An individual's encoding (green circles) and recognition (purple squares) fixations on a given 360 image. The bright green (solid) insert and Panel B show an example of a deterministic sequence. The pink inset and Panel C show an example of vertical laminarity. This participant came back to inspect this region in more detail during recognition. Panel D is the cross recurrence plot resulting from the overlap of these encoding and recognition fixation sequences.

Figure 3 depicts the participant means for each CRA measure for correctly recognized images and those that were forgotten. In order to determine whether a given CRA measure predicted memory performance, each measure was submitted to a GLMM. Including participants and images as random effects, this GLMM takes into account a given person's ability to recognize images, as well as a given image's memorability (details on p. 8). Table 2 shows the likelihood ratio tests and model estimates output from the GLMM for each CRA measure for gaze and head sequences.

Figure 3: Participant means (in colours) and overall mean (black) cross recurrence measures for gaze and head comparisons across correct and incorrectly recognized trials.

Table 2: Likelihood ratio tests and model estimates for whether each CRA measure predicts memory performance (significant predictors bolded).

Measure	χ^2	<i>p</i>	<i>Odds Ratio</i>	<i>CI</i>	<i>z</i>	<i>p</i>
Gaze						
Recurrence	2.74	.098	1.26	[0.94, 1.70]	1.53	.127
Determinism	1.99	.158	1.18	[0.94, 1.48]	1.44	.150
Maxline	1.37	.241	1.15	[0.91, 1.45]	1.18	.238
Vertical Laminarity	6.04	.014	1.31	[1.06, 1.62]	2.54	.011
Vertical Trapping Time	26.93	< .001	2.43	[1.64, 3.59]	4.45	< .001
Horizontal Laminarity	< 1		1.02	[0.82, 1.28]	< 1	
Horizontal Trapping Time	< 1		0.96	[0.77, 1.20]	< 1	
Corm	5.55	.018	1.31	[1.05, 1.64]	2.39	.017
Clusters	4.19	.041	1.27	[1.01, 1.60]	2.04	.041
Head						
Recurrence	< 1		1.05	[0.83, 1.32]	< 1	
Determinism	< 1		1.16	[0.93, 1.44]	1.33	.180
Maxline	1.19	.275	1.14	[0.90, 1.43]	1.09	.277
Vertical Laminarity	13.10	< .001	1.47	[1.20, 1.81]	3.67	< .001
Vertical Trapping Time	16.53	< .001	2.27	[1.38, 3.75]	3.21	.001
Horizontal Laminarity	< 1		0.98	[0.79, 1.23]	< 1	
Horizontal Trapping Time	4.23	.040	0.82	[0.69, 0.98]	2.23	.026
Corm	4.96	.026	1.29	[1.03, 1.61]	2.26	.024
Clusters	< 1		1.07	[0.87, 1.32]	< 1	

For gaze, Vertical Laminarity and Trapping Time, as well as CORM and the Clusters measure significantly predicted memory performance. Higher Vertical Laminarity, Trapping Time and Cluster percentages for correctly recognized images indicates that successful recognition depended on fixating several previously encoded regions. Higher CORM indicates that a given image was more likely to be recognized when participants re-fixated regions that were encoded earlier in the trial – a primacy effect.

For head positions, Vertical Laminarity and Trapping Time, as well as Horizontal Trapping Time and CORM significantly predicted memory performance. Similar to gaze, higher Vertical Laminarity and Trapping Time for correctly recognized images reflects head movements during recognition that overlap with those made during encoding. Interestingly, correct recognition was also associated with reduced Horizontal Trapping Time (along with no effect of Horizontal Laminarity), indicating that observers explore regions during recognition that were explored during encoding, but quickly move on to new regions (see Figure 4 for an example of a trial and cross recurrence plot with particularly low Horizontal Trapping Time). This aligns well with the idea that the greater exploration (i.e., moving the head to unseen regions) during encoding and recognition improve memory performance. Higher CORM indicates a primacy effect, where it is particularly important for recognition that the overlap occurs at the start of the trial.

In CRA, many of the measures are highly correlated. In order to take this into account, and determine whether a given similarity measure predicted performance over and above the others, we included each of the significant measures from above into an overall GLMM, with each measure as a fixed effect. This was done separately for the gaze and head data. For gaze, Vertical Trapping Time predicted memory performance, over and above the other measures, *Odds Ratio* =

2.33, $CI = [1.50, 3.64]$, $z = 3.73$, $p < .001$, and a model including the Vertical Trapping Time measure significantly outperformed a model without it, $\chi^2(1) = 17.36$, $p < .001$.

For head positions, the fixed effects Vertical Laminarity, $\chi^2(1) = 4.56$, $p = .033$, *Odds Ratio* = 1.31, $CI = [1.03, 1.67]$, $z = 2.17$, $p = .030$, Vertical Trapping Time, $\chi^2(1) = 9.45$, $p = .002$, *Odds Ratio* = 2.02, $CI = [1.18, 3.44]$, $z = 2.57$, $p = .010$, and Horizontal Trapping Time, $\chi^2(1) = 10.74$, $p = .001$, *Odds Ratio* = 0.68, $CI = [0.55, 0.85]$, $z = -3.44$, $p < .001$, all uniquely predicted memory performance, with the strongest predictor being Horizontal Trapping Time. Horizontal Trapping Time was significantly smaller in correctly recognized images, compared to those that were not, possibly reflecting that once an image was correctly recognized observers felt free to move on from regions examined during encoding (see Figure 4).

Figure 4: Example trial with low Horizontal Trapping Time. Blue circles are encoding head positions, while orange squares are recognition head positions. Note how head positions during recognition hit several regions that were previously encoded.

In summary, successful recognition of a panoramic scene appears to be characterized by increased vertical laminarity and trapping time for both gaze and head positions, indicating that recognition is associated with reorienting (with the eyes and head) to previously encoded regions. Interestingly, vertical trapping time values of approximately 3 fixations and 6 head positions are made at previously encoded regions (see Figure 3). For head movements, successful recognition was most strongly predicted by horizontal trapping time, which is in line with greater exploration with the head during correctly recognized trials. In addition, it appears that it is particularly important that fixation and head positions overlap early on during encoding and recognition trials - a primacy effect.

3.8. ROC analysis of CRA measures

In order to determine if including CRA measures of head movements leads to a stronger connection between motor movements and subsequent memory performance, each of the gaze CRA measures, each of the head CRA measures, and the combination of both were used to predict correct trials via separate logistic regression models (i.e., a gaze-only model, a head-only model, and a combined model using all the CRA measures as predictors). Subjects and images were included as fixed effects as in the previous GLMMs. These models were plotted as receiver operating characteristic (ROC) curves because they allow for the evaluation of a collection of continuous markers to predict a binary outcome—in this case, recognition (Robin et al., 2011; Seshan et al., 2013). They enable the direct comparison of a new marker (i.e., predictor) on a model's overall predictive ability in an intuitive way. Thus, we can directly evaluate the predictive power associated with using CRA head measures in addition to CRA gaze measures. To compare two ROC curves, we computed their areas under the curve (AUC). The difference between AUC_i and AUC_j was determined using $D = (AUC_i - AUC_j) / sd(AUC_i - AUC_j)$ (see Robin et al., 2011),

where $sd(AUC_i - AUC_j)$ was estimated using bootstrapping with 2000 bootstrap replications (Carpenter & Bithell, 2000). The classification performance of these three models, and their corresponding 95% CI's is displayed in Figure 5.

Figure 5: Classification performance of the gaze CRA measures (green), the head CRA measures (orange), and a model combining the two (purple). 95% CI's given for each AUC in brackets.

All three GLMM models performed significantly above chance reflecting the association between motor movements and subsequent memory performance. The model containing the CRA measures for the head, did not significantly outperform the model containing the CRA measures for gaze, $D = 1.17$, $p = .24$, nor did it outperform the combined head and gaze model, $D = 0.93$, $p = .35$. Most notably, the combined model containing CRA measures for the head and gaze was more accurate than the model containing CRA measures for gaze movements alone, $D = 2.17$, $p = .03$. This indicated that the clearest associations between motor movements and memory performance were when gaze and head movements were taken into account from both encoding and recognition phases. This supports our hypothesis that the strongest evidence for Scanpath Theory using CRA measures will emerge when both head and eye measures are used to predict recognition, as opposed to eye movements alone.

4. Discussion

In the present experiment, we tested Scanpath Theory in 360° scenes presented in a VR headset. Participants were asked to encode 80 scenes for 10s each, and then recognize a subset of the encoded scenes. We analyzed the spatial and temporal overlap of unconstrained head and eye movements when encoding and recognizing these scenes, and assessed the extent to which the overlap predicted recognition performance. We were interested in two main questions, 1) whether certain spatial and temporal aspects eye and head movement behaviours predicted recognition performance, and 2) whether the inclusion of head movement measures increased the ability to predict recognition performance – strengthening the functional role of motor movements in Noton and Stark's (1971b) Scanpath Theory. We consider each question in turn below.

4.1. Do spatial and temporal aspects of eye and head movements predict correct recognition?

We found that many spatial and temporal gaze movement patterns predicted correct recognition. Scenes that were later correctly recognized were explored more thoroughly by participants during encoding. When recognizing a scene, observers made detailed re-inspections, looking at one or a few regions that were previously encoded, making on average 3 fixations at each of these regions. In addition, it was particularly important that the overlap of fixation positions at encoding and recognition happened early on (a primacy effect).

A similar pattern was observed for head movements. Observers were more likely to correctly recognize scenes when they moved their head to explore it more thoroughly during encoding. In addition, greater overlap of head positions at encoding and recognition across a few scene regions led to better recognition performance, and this overlap was more likely to occur during early scene viewing (a primacy effect). Unlike gaze, however, for head movements the

strongest predictor of recognition performance was a spatio-temporal measure that captured continuous exploratory head movements during recognition, where the head moved to overlap with previously examined regions without dwelling too long at any one. Specifically, the head rarely remained in any one location for longer than it takes to make four consecutive fixations (see Figure 3).

Collectively, these results suggest that while scanpaths and headpaths in panoramic scenes are not repeated one-to-one (i.e., a strict reading of Scanpath Theory is not upheld), certain eye and head movement behaviours, aside from exploration, lead to better recognition performance (and presumably, better scene encoding). During 360° scene recognition, observers move their eyes and head to a few locations that were previously encoded. In a prototypical situation, when a previously viewed scene is correctly recognized, the head tends to cover regions that were previously visited, presumably so that the eyes can examine key regions. Critically, the overlap of gaze and head positions show a primacy effect, where these revisits tend to occur earlier in time. This suggests that recognition occurs rapidly, within the first few fixations. These results paint an interesting and detailed picture of how observers use their eye and head movements to encode and recognize 360° scenes.

4.2. Does the inclusion of head movement measures strengthen the functional role of motor movements?

A second question we sought to address was whether a general version of Scanpath Theory garners greater support when observers are free to move their head and eyes to encode and recognise a scene. To answer this question, we examined recognition performance as a function of (a) spatio-temporal gaze measures, (b) spatio-temporal head measures, and (c) their combination.

We found that a combined model, including all the gaze and head measures, outperformed a model that contained only the gaze measures in predicting recognition performance. The equivalence of head measures alone, and head and gaze, speak to the fact that head movements tend to function in service of the eyes (e.g., Solman et al., 2017). That is, the head moves so that the eyes can take in the visual world. Collectively these findings indicate that including head movement information in addition to gaze data strengthens the evidence for Scanpath Theory.

This work aligns well with ongoing calls from proponents of embodied cognition to avoid using overly reductionistic methods because they sacrifice the complexity of ongoing interactions between the brain, body, and environment (e.g., Núñez, 2012). The present findings specifically fit into this broader narrative as our results directly revealed that had we taken a more traditional, reductionistic approach (i.e., by excluding head movements), the association would have been less clear between participants' motor activity and their memory. A notable benefit of using VR was that there was no sacrifice in experimental control as tradeoff for taking the more holistic, embodied approach. Ironically, the stimuli used in the present experiment were more realistic than prototypical stimuli as they were completely immersive scenes unlike previous Scanpath Theory investigations that typically use computer monitors. This highlights VR as a methodological tool that allows for controlled, yet embodied approaches to investigating action and cognition.

5. Conclusions

Our study suggests that certain repeated eye and head movement behaviours during recognition are functional for memory. Rather than reinstating scan-, or head-paths one-to-one, observers fixate to a few key diagnostic regions that were encoded, with the head operating in service of the eyes so that these regions can be visually inspected. Crucially, it was revealed that

when incorporating both head and gaze measures, the clearest associations between visuomotor behaviour and memory emerged. This suggests that head movements are supportive of successful scene recognition and by extension are a critical part of the visual system. Furthermore, since this was the case for head CRA measures in particular, it suggests that the head not only plays an important role in attentional allocation by enabling the eyes to fixate on regions of interest; rather, the head plays a vital role in ongoing scene perception and recognition. While Noton's "interesting but rather informally presented theory" (Allen & Noton, 1971) may not have received unequivocal support in the past, evidence that favours a more general version of this theory is discovered when the dynamics of eye and head movements are considered within a (virtually) real panoramic environment.

Authorship Contributions

Nicola C. Anderson: conceptualization, methodology, investigation, data analysis, writing – original draft, review and editing. Oliver Jacobs: conceptualization, data analysis, writing – review and editing. Walter F. Bischof: conceptualization, methodology, investigation, data analysis, writing – review and editing. Alan Kingstone: conceptualization, methodology, writing – review and editing

Declaration of Competing Interest

The authors declare that they have no known competing financial interests of personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research was supported by the Natural Sciences and Engineering Research Council of Canada (NCA: Postdoctoral Fellowship; WFB: RGPIN-2015-05560; AK: RGPIN-2016-04319).

6. References

- Allen, A. D., & Noton, D. (1971). Comments on “A Theory of Visual Pattern Perception.” *IEEE Transactions on Systems, Man, and Cybernetics*, *SMC-1*(4), 388–389.
<https://doi.org/10.1109/TSMC.1971.4308325>
- Althoff, R. R., & Cohen, N. J. (1999). Eye-movement-based memory effect: A reprocessing effect in face perception. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(4), 997.
- Anderson, N. C., Anderson, F., Kingstone, A., & Bischof, W. F. (2015). A comparison of scanpath comparison methods. *Behavior Research Methods*, *47*(4), 1377–1392.
<https://doi.org/10.3758/s13428-014-0550-3>
- Anderson, N. C., Bischof, W. F., Laidlaw, K. E. W., Risko, E. F., & Kingstone, A. (2013). Recurrence quantification analysis of eye movements. *Behavior Research Methods*, *45*(3), 842–856. <https://doi.org/10.3758/s13428-012-0299-5>
- Backhaus, D., Engbert, R., Rothkegel, L. O. M., & Trukenbrod, H. A. (2020). Task-dependence in scene perception: Head unrestrained viewing using mobile eye-tracking. *Journal of Vision*, *20*(5), 3–3. <https://doi.org/10.1167/jov.20.5.3>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.
<https://doi.org/10.1016/j.jml.2012.11.001>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1), 1–48.
<https://doi.org/10.18637/jss.v067.i01>

- Bischof, W. F., Anderson, N. C., Doswell, M. T., & Kingstone, A. (2020). *Visual Exploration of Omni-Directional Panoramic Scenes*. PsyArXiv. <https://doi.org/10.31234/osf.io/vpxm4>
- Bischof, W. F., Anderson, N. C., & Kingstone, A. (2019). Temporal Methods for Eye Movement Analysis. In C. Klein & U. Ettinger (Eds.), *Eye Movement Research: An Introduction to its Scientific Foundations and Applications* (pp. 407–448). Springer International Publishing. https://doi.org/10.1007/978-3-030-20085-5_10
- Brandt, S. A., & Stark, L. W. (1997). Spontaneous eye movements during visual imagery reflect the content of the visual scene. *Journal of Cognitive Neuroscience*, 9(1), 27–38.
- Buswell, G. T. (1935). *How people look at pictures*. University of Chicago Press Chicago.
- Carpenter, J., & Bithell, J. (2000). Bootstrap confidence intervals: When, which, what? A practical guide for medical statisticians. *Statistics in Medicine*, 19(9), 1141–1164.
- Cristino, F., Mathôt, S., Theeuwes, J., & Gilchrist, I. D. (2010). ScanMatch: A novel method for comparing fixation sequences. *Behavior Research Methods*, 42(3), 692–700. <https://doi.org/10.3758/BRM.42.3.692>
- Dale, R., Warlaumont, A. S., & Richardson, D. C. (2011). Nominal cross recurrence as a generalized lag sequential analysis for behavioral streams. *International Journal of Bifurcation and Chaos*, 21(04), 1153–1161.
- Damiano, C., & Walther, D. B. (2019). Distinct roles of eye movements during memory encoding and retrieval. *Cognition*, 184, 119–129. <https://doi.org/10.1016/j.cognition.2018.12.014>
- Dewhurst, R., Nyström, M., Jarodzka, H., Foulsham, T., Johansson, R., & Holmqvist, K. (2012). It depends on how you look at it: Scanpath comparison in multiple dimensions with MultiMatch, a vector-based approach. *Behavior Research Methods*, 44(4), 1079–1100. <https://doi.org/10.3758/s13428-012-0212-2>

- Fahimi, R., & Bruce, N. D. (2020). On metrics for measuring scanpath similarity. *Behavior Research Methods*, 1–20.
- Foulsham, T., Dewhurst, R., Nyström, M., Jarodzka, H., Johansson, R., Underwood, G., & Holmqvist, K. (2012). *Comparing scanpaths during scene encoding and recognition: A multi-dimensional approach*. <http://dspace.ou.nl/handle/1820/4627>
- Foulsham, T., & Kingstone, A. (2013). Fixation-dependent memory for natural scenes: An experimental test of scanpath theory. *Journal of Experimental Psychology: General*, 142(1), 41–56. <https://doi.org/10.1037/a0028227>
- Foulsham, T., & Underwood, G. (2008). What can saliency models predict about eye movements? Spatial and sequential aspects of fixations during encoding and recognition. *Journal of Vision*, 8(2), 6–6. <https://doi.org/10.1167/8.2.6>
- Foulsham, T., Walker, E., & Kingstone, A. (2011). The where, what and when of gaze allocation in the lab and the natural environment. *Vision Research*, 51(17), 1920–1931. <https://doi.org/10.1016/j.visres.2011.07.002>
- Godijn, R., & Theeuwes, J. (2012). Overt is no better than covert when rehearsing visuo-spatial information in working memory. *Memory & Cognition*, 40(1), 52–61.
- Harding, G., & Bloj, M. (2010). Real and predicted influence of image manipulations on eye movements during scene recognition. *Journal of Vision*, 10(2), 8–8.
- Hessels, R. S., Niehorster, D. C., Nyström, M., Andersson, R., & Hooge, I. T. (2018). Is the eye-movement field confused about fixations and saccades? A survey among 124 researchers. *Royal Society Open Science*, 5(8), 180502.

- Hollingworth, A., & Henderson, J. M. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance*, 28(1), 113–136. <https://doi.org/10.1037/0096-1523.28.1.113>
- Holm, L., & Mäntylä, T. (2007). Memory for scenes: Refixations reflect retrieval. *Memory & Cognition*, 35(7), 1664–1674.
- Hooge, I. T., Hessels, R. S., Niehorster, D. C., Diaz, G. J., Duchowski, A. T., & Pelz, J. B. (2019). *From lab-based studies to eye-tracking in virtual and real worlds: Conceptual and methodological problems and solutions.*
- Humphrey, K., & Underwood, G. (2010). The potency of people in pictures: Evidence from sequences of eye fixations. *Journal of Vision*, 10(10), 19–19.
- Johansson, R., Holsanova, J., Dewhurst, R., & Holmqvist, K. (2011). Eye movements during pictorial recall have a functional role, but they are not reinstatements of those from encoding. *Journal of Eye Movement Research*, 4(3), 32–32.
- Kolers, P. A. (1973). Remembering operations. *Memory & Cognition*, 1(3), 347–355.
- Laeng, B., & Teodorescu, D.-S. (2002). Eye scanpaths during visual imagery reenact those of perception of the same visual scene. *Cognitive Science*, 26(2), 207–231.
- Le Meur, O., & Baccino, T. (2013). Methods for comparing scanpaths and saliency maps: Strengths and weaknesses. *Behavior Research Methods*, 45(1), 251–266. <https://doi.org/10.3758/s13428-012-0226-9>
- Lenth, R. (2019). *emmeans: Estimated Marginal Means, aka Least-Squares Means.* <https://CRAN.R-project.org/package=emmeans>

- Li, C.-L., Aivar, M. P., Kit, D. M., Tong, M. H., & Hayhoe, M. M. (2016). Memory and visual search in naturalistic 2D and 3D environments. *Journal of Vision*, 16(8), 9–9. <https://doi.org/10.1167/16.8.9>
- Li, C.-L., Aivar, M. P., Tong, M. H., & Hayhoe, M. M. (2018). Memory shapes visual search strategies in large-scale environments. *Scientific Reports*, 8(1), 4324. <https://doi.org/10.1038/s41598-018-22731-w>
- Locher, P. J., & Nodine, C. F. (1974). The role of scanpaths in the recognition of random shapes. *Perception & Psychophysics*, 15(2), 308–314.
- Lüdecke, D. (2020). *sjPlot: Data Visualization for Statistics in Social Science*. <https://CRAN.R-project.org/package=sjPlot>
- Mannan, S., Ruddock, K. H., & Wooding, D. S. (1995). Automatic control of saccadic eye movements made in visual inspection of briefly presented 2-D images. *Spatial Vision*.
- Noton, D. (1970). A theory of visual pattern perception. *IEEE Transactions on Systems Science and Cybernetics*, 6(4), 349–357.
- Noton, D., & Stark, L. (1971a). Eye movements and visual perception. *Scientific American*, 224(6), 34–43.
- Noton, D., & Stark, L. (1971b). Scanpaths in eye movements during pattern perception. *Science*, 171(3968), 308–311.
- Noton, D., & Stark, L. (1971c). Scanpaths in saccadic eye movements while viewing and recognizing patterns. *Vision Research*, 11(9), 929-IN8. [https://doi.org/10.1016/0042-6989\(71\)90213-6](https://doi.org/10.1016/0042-6989(71)90213-6)

- Núñez, R. (2012). On the Science of Embodied Cognition in the 2010s: Research Questions, Appropriate Reductionism, and Testable Explanations. *Journal of the Learning Sciences*, 21(2), 324–336. <https://doi.org/10.1080/10508406.2011.614325>
- Pixel connectivity. (2021). In *Wikipedia*. https://en.wikipedia.org/wiki/Pixel_connectivity
- Rajsic, J., Wilson, D. E., & Pratt, J. (2018). The price of information: Increased inspection costs reduce the confirmation bias in visual search. *Quarterly Journal of Experimental Psychology*, 71(4), 832–849. <https://doi.org/10.1080/17470218.2016.1278249>
- Richardson, D. C., Altmann, G. T., Spivey, M. J., & Hoover, M. A. (2009). Much ado about eye movements to nothing: A response to Ferreira et al.: Taking a new look at looking at nothing. *Trends in Cognitive Sciences*, 13(6), 235–236.
- Richardson, D. C., & Spivey, M. J. (2000). Representation, space and Hollywood Squares: Looking at things that aren't there anymore. *Cognition*, 76(3), 269–295.
- Risko, E. F., Medimorec, S., Chisholm, J., & Kingstone, A. (2014). Rotating with rotated text: A natural behavior approach to investigating cognitive offloading. *Cognitive Science*, 38(3), 537–564.
- Risko, E. F., Richardson, D. C., & Kingstone, A. (2016). Breaking the fourth wall of cognitive science: Real-world social attention and the dual function of gaze. *Current Directions in Psychological Science*, 25(1), 70–74.
- Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., & Müller, M. (2011). PROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics*, 12, 77.
- Seshan, V. E., Gönen, M., & Begg, C. B. (2013). Comparing ROC curves derived from regression models. *Statistics in Medicine*, 32(9), 1483–1493. <https://doi.org/10.1002/sim.5648>

- Singmann, H., Bolker, B., Westfall, J., Aust, F., & Ben-Shachar, M. S. (2019). *afex: Analysis of Factorial Experiments*. <https://CRAN.R-project.org/package=afex>
- Solman, G. J., Foulsham, T., & Kingstone, A. (2017). Eye and head movements are complementary in visual selection. *Royal Society Open Science*, *4*(1), 160569.
- Solman, G. J., & Kingstone, A. (2014). Balancing energetic and cognitive resources: Memory use during search depends on the orienting effector. *Cognition*, *132*(3), 443–454. <https://doi.org/10.1016/j.cognition.2014.05.005>
- 't Hart, B. M., Vockeroth, J., Schumann, F., Bartl, K., Schneider, E., König, P., & Einhäuser, W. (2009). Gaze allocation in natural stimuli: Comparing free exploration to head-fixed viewing conditions. *Visual Cognition*, *17*(6–7), 1132–1158. <https://doi.org/10.1080/13506280902812304>
- Tremblay, S., Saint-Aubin, J., & Jalbert, A. (2006). Rehearsal in serial memory for visual-spatial information: Evidence from eye movements. *Psychonomic Bulletin & Review*, *13*(3), 452–457.
- Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, *80*(5), 352.
- Unity Technologies (2017.4.1). (2017). [Computer software]. Unity Technologies. unity3d.com
- Võ, M. L.-H., & Wolfe, J. M. (2015). The role of memory for visual search in scenes. *Annals of the New York Academy of Sciences*, *1339*(1), 72–81.
- Webber Jr, C. L., & Zbilut, J. P. (2005). Recurrence quantification analysis of nonlinear dynamical systems. *Tutorials in Contemporary Nonlinear Methods for the Behavioral Sciences*, *94*(2005), 26–94.

- Whiteside, J. A. (1978). Visual display recognition and the duplication of inspection sequences. *Perception & Psychophysics*, 23(4), 350–355.
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York.
<https://ggplot2.tidyverse.org>
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., & Torralba, A. (2010). Sun database: Large-scale scene recognition from abbey to zoo. *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference On*, 3485–3492.
<http://ieeexplore.ieee.org/abstract/document/5539970/>
- Yarbus, A. L. (1967). *Eye movements during perception of complex objects*. Springer.
http://link.springer.com/10.1007/978-1-4899-5379-7_8
- Zelinsky, G. J., & Loschky, L. C. (2005). Eye movements serialize memory for objects in scenes. *Perception & Psychophysics*, 67(4), 676–690.

Appendix 1: Cross-Recurrence Analysis

Consider two fixation sequences $\mathbf{e} = \mathbf{e}_i, i=1, \dots, N$, with $\mathbf{e}_i = \langle x_i, y_i \rangle$ and $\mathbf{r} = \mathbf{r}_j, j=1, \dots, N$, with $\mathbf{r}_j = \langle x_j, y_j \rangle$. For fixation sequences of unequal length, the longer sequence is truncated. Two fixations $\mathbf{e}_i$ and $\mathbf{r}_j$ are cross-recurrent if they are close together, i.e., we define the cross-recurrence of two fixations c_{ij} as

$$c_{ij} = \begin{cases} 1, & d(\mathbf{e}_i, \mathbf{r}_j) \leq \rho \\ 0, & \text{otherwise} \end{cases}$$

where d is a distance measure (e.g., Euclidian distance in 2D or the great circle (angular) distance in spherical displays) and ρ is a threshold.

Several measures can be used for characterizing cross-recurrence patterns. The measures are extensions of the recurrence measures introduced by Anderson et al. (2013). In the cross-recurrence matrix $\mathbf{C} = c_{ij}, i=1, \dots, N, j=1, \dots, N$; let D be a diagonal line representing a fixation trajectory common to both fixation sequences; let L_H be a horizontal laminate line, representing a single fixation in $\mathbf{r}$ being recurrent with a sequence of fixations in $\mathbf{e}$; and let L_V be vertical laminate line, representing a single fixation in $\mathbf{e}$ being recurrent with a sequence of fixations in $\mathbf{r}$. D, L_H and L_V are taken into account only if their lengths are at least L_{min} (usually $L_{min}=2$). Finally, let C be the sum of recurrences, i.e., $C = \sum_{i=1}^N \sum_{j=1}^N c_{ij}$ and let $|\cdot|$ denote cardinality.

Cross-Recurrence. The Cross-Recurrence measure of two fixation sequences defined as

$$REC = 100 \cdot \frac{C}{N^2}.$$

It represents the percentage of cross-recurrent fixations, i.e., the percentage of fixations that match (i.e., are close) between the two fixation sequences.

Determinism. The Determinism measure is defined as

$$DET = 100 \cdot \frac{|D|}{C}.$$

It is the percentage of cross-recurrent points that form diagonal lines and represents the percentage of fixation trajectories common to both fixation sequences. That is, it quantifies the overlap of a specific sequence of fixations, preserving the sequential information.

Vertical Laminarity. The horizontal laminarity measure is defined as

$$vLAM = 100 \cdot \frac{|L_V|}{C}.$$

It is the proportion of fixations in vertical laminate lines L_V and represents the percentage of fixation that were fixated in detail in $\mathbf{r}$, but only fixated briefly in $\mathbf{e}$.

Horizontal Laminarity. The horizontal laminarity measure is defined as

$$hLAM = 100 \cdot \frac{|L_H|}{C}.$$

It is the proportion of fixations in horizontal laminate lines L_H and represents the percentage of fixation that were fixated in detail in $\mathbf{e}$, but only fixated briefly in $\mathbf{r}$.

Vertical Trapping Time. The vertical trapping time vTT is defined as the average length of all vertical laminate lines.

Horizontal Trapping Time. The horizontal trapping time hTT is defined as the average length of all horizontal laminate lines.

Clusters. A cluster in the cross-recurrence matrix is defined by the following connectivity criterion: If the cross-recurrence c is in a cluster, then the cross-recurrences in the 8-connected neighbourhood (“Pixel Connectivity,” 2021) of c are also in the cluster. Only clusters with at least 8 cross-recurrences are counted. The *CLUST* measure is defined as the percentage of cross-recurrent fixations that are part of a cluster.

Center of Recurrence Mass. The center of recurrence mass *corm* is defined as the distance of the center of gravity from the main diagonal, normalized such that the maximum possible value is 100,

$$CORM = 100 \frac{\sum_{i=1}^N \sum_{j=1}^N (j-i) r_{ij}}{(N-1)C}.$$

The *corm* measure indicates the dominant lag of cross-recurrences. Small *corm* values indicate that the same fixations in both fixation sequences tend to occur close in time, whereas large *corm* values indicate that cross-recurrences tend to occur with either a large positive or negative lag. *CORM* values are positive if the fixation sequence $\mathbf{e}$ is leading the fixation sequence $\mathbf{r}$, and it is negative if the fixation sequence $\mathbf{r}$ is leading the fixation sequence $\mathbf{e}$.